

**NAME**

**recoverdisk** - recover data from hard disk or optical media

**SYNOPSIS**

**recoverdisk** [-b *bigsize*] [-r *readlist*] [-s *interval*] [-u *pattern*] [-v] [-w *writelist*] *source* [*destination*]

**DESCRIPTION**

The **recoverdisk** utility reads data from the *source* file until all blocks could be successfully read. If *destination* was specified all data is being written to that file. It starts reading in multiples of the sector size. Whenever a block fails, it is put to the end of the working queue and will be read again, possibly with a smaller read size.

By default it uses block sizes of roughly 1 MB, 32kB, and the native sector size (usually 512 bytes). These figures are adjusted slightly, for devices whose sectorsize is not a power of 2, e.g., audio CDs with a sector size of 2352 bytes.

The options are as follows:

**-b** *bigsize*

The size of reads attempted first. The middle pass is roughly the logarithmic average of the *bigsize* and the sectorsize.

**-r** *readlist*

Read the list of blocks and block sizes to read from the specified file.

**-s** *interval*

How often we should update the writelist file while things go OK. The default is 60 and the unit is "progress messages" so if things go well, this is the same as once per minute.

**-u** *pattern*

By default blocks which encounter read errors will be filled with the pattern ‘\_UNREAD\_’ in the output file. This option can be used to specify another pattern. Nothing gets written if the string is empty.

**-v** Enables nicer status report using ANSI escapes and UTF-8.

**-w** *writelist*

Write the list of remaining blocks to read to the specified file if **recoverdisk** is aborted via SIGINT.

The **-r** and **-w** options can be specified together. Especially, they can point to the same file, which will be updated on abort.

## OUTPUT

The **recoverdisk** utility prints several columns, detailing the progress

*start* Starting offset of the current block.

*size* Read size of the current block.

*len* Length of the current block.

*state* Is increased for every failed read.

*done* Number of bytes already read.

*remaining* Number of bytes remaining.

*% done* Percent complete.

## EXAMPLES

```
# recover data from failing hard drive ada3
recoverdisk /dev/ada3 /data/disk.img
```

```
# clone a hard disk
recoverdisk /dev/ada3 /dev/ada4
```

```
# read an ISO image from a CD-ROM
recoverdisk /dev/cd0 /data/cd.iso
```

```
# continue reading from a broken CD and update the existing worklist
recoverdisk -r worklist -w worklist /dev/cd0 /data/cd.iso
```

```
# recover a single file from the unreadable media
recoverdisk /cdrom/file.avi file.avi
```

```
# If the disk hangs the system on read-errors try:
recoverdisk -b 0 /dev/ada3 /somewhere
```

**SEE ALSO**

dd(1), ada(4), cam(4), cd(4), da(4)

**HISTORY**

The **recoverdisk** utility first appeared in FreeBSD 7.0.

**AUTHORS**

The original implementation was done by Poul-Henning Kamp <[phk@FreeBSD.org](mailto:phk@FreeBSD.org)> with minor improvements from Ulrich Spörlein <[uqs@FreeBSD.org](mailto:uqs@FreeBSD.org)>.

This manual page was written by Ulrich Spörlein.

**BUGS**

Reading from media where the sectorsize is not a power of 2 will make all 1 MB reads fail. This is due to the DMA reads being split up into blocks of at most 128kB. These reads then fail if the sectorsize is not a divisor of 128kB. When reading a full raw audio CD, this leads to roughly 700 error messages flying by. This is harmless and can be avoided by setting **-b** to no more than 128kB.

**recoverdisk** needs to know about read errors as fast as possible, i.e., retries by lower layers will usually slow down the operation. When using cam(4) attached drives, you may want to set kern.cam.XX.retry\_count to zero, e.g.:

```
# sysctl kern.cam.ada.retry_count=0
# sysctl kern.cam.cd.retry_count=0
# sysctl kern.cam.da.retry_count=0
```